arXiv:1409.2702v1 [cs.CV] 9 Sep 2014

F-formation Detection: Individuating Free-standing Conversational Groups in Images

Francesco Setti*[1], Chris Russell[2], Chiara Bassetti[1], and Marco Cristani[3,4]

[1]Institute of Cognitive Science and Technologies (ISTC), Italian National Research Council (CNR), Trento, Italy
[2]Department of Computer Science, University College London, UK
[3]Department of Computer Science, University of Verona, Italy
[4]Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia Genova(IIT), Italy

Abstract

Detection of groups of interacting people is a very interesting and useful task in many modern technologies, with application fields spanning from video-surveillance to social robotics. In this paper we first furnish a rigorous definition of group considering the background of the social sciences: this allows us to specify many kinds of group, so far neglected in the Computer Vision literature. On top of this taxonomy, we present a detailed state of the art on the group detection algorithms. Then, as a main contribution, we present a brand new method for the automatic detection of groups in still images, which is based on a graph-cuts framework for clustering individuals; in particular we are able to codify in a computational sense the sociological definition of F-formation, that is very useful to encode a group having only proxemic information: position and orientation of people. We call the proposed method *Graph-Cuts for F-formation* (GCFF). We show how GCFF definitely outperforms all the state of the art methods in terms of different accuracy measures (some of them are brand new), demonstrating also a strong robustness to noise and versatility in recognizing groups of various cardinality.

1 Introduction

After years of research on automated analysis of individuals, the computer vision community has transferred its attention on the new issue of modeling gatherings of people, commonly referred as *groups* [23, 17, 16, 24].

*francesco.setti@loa.istc.cnr.it

A group can be broadly understood as a social unit comprising several members who stand in status and relationships with one another [19]. However, there are many kinds of groups, that differ in dimension (small groups or crowds), durability (ephemeral, *ad hoc* or stable groups), in/formality of organization, degree of "sense of belonging", level of physical dispersion etc. [25] (see the literature review in the next section). In this article, we build from the concepts of sociological analysis and we focus on free-standing conversational groups (FCGs), or small ensembles of co-present persons engaged in *ad hoc* focused encounters [25, 26, 38]. FCGs represent crucial social situations, and one of the most fundamental bases of dynamic sociality: these facts make them a crucial target for the modern automated monitoring and profiling strategies which have started to appear in the literature in the last three years [16, 6, 31, 57, 58, 62].

In computer vision, the analysis of groups has occurred historically in two broad contexts: video-surveillance and meeting analysis.

Within the scope of video-surveillance, the definition of a group is generally simplified to two or more people of similar velocity, spatially and temporally close to one another [22]. This simplified definition arises from the difficulty of inferring persistent social structure from short video clips. In this case, most of the vision-based approaches perform group tracking, i.e. capturing individuals in movement and maintaining their identity across video frames, understanding how they are partitioned in groups [44, 24, 53, 7, 48, 22].

In meeting analysis, typified by classroom behavior [23], people typically sit around a table and remain near a fixed location for most of the time, predominantly interacting through speech and gesture. In such a scenario, activities can be finely monitored using a variety of audiovisual features, captured by pervasive sensors like portable devices, microphone arrays, etc. [66, 34, 30].

From a sociological point of view, meetings are examples of social organization that employs *focused* interaction, which occurs when persons openly cooperate to sustain a single focus of attention [25, 26]. This broad definition covers other collaborative situated systems of activity that entail a more or less static spatial and proxemic organization – such as playing a board or sport game, having dinner, doing a puzzle together, pitching a tent, or free conversation [25], whether sitting on the couch at a friend's place, standing in the foyer and discussing the movie, or leaning on the balcony and smoking a cigarette during work-break.

Free-standing conversational groups (FCGs) [38] are another example of *focused* encounters. FCGs emerge during many and diverse social occasions, such as a party, a social dinner, a coffee break, a visit in a museum, a day at the seaside, a walk in the city plaza or at the mall; more generally, when people spontaneously decide to be in each other's immediate presence to interact with one another. For these reasons, FCGs are fundamental social entities, whose automated analysis may bring to a novel level of activity and behavior analysis.

In a FCG, people communicate to the other participants, among –and above all– the rest, what they think they are doing together, what they regard as the activity at hand. And they do so not only, and perhaps not so much, by talking, but also, and as much, by exploiting non-verbal modalities of expression, also

called social signals [63], among which positional and orientational forms play a crucial role (cf. also [26], p. 11). In fact, the spatial position and orientation of people define one of the most important proxemic notions which describe an FCG, that is, the Adam Kendon's *Facing Formation*, mostly known as *F-formation*.

In Kendon's terms [39, 14, 38], an F-formation is a socio-spatial formation in which people have established and maintain a convex space (called *o-space*) to which everybody in the gathering has direct, easy and equal access. Typically, people arrange themselves in a form of a circle, ellipse, horseshoe, side-by-side or L-shape (cf. Fig. 4), so that they can have easy and preferential access to one another while excluding distractions of the outside world with their backs. Examples of F- formations are reported in Fig. 1. In computer vision, spatial position and orientational information can be automatically extracted, and these facts pave the way to the computational modeling of F-formation and, as a consequence, of the FCGs.

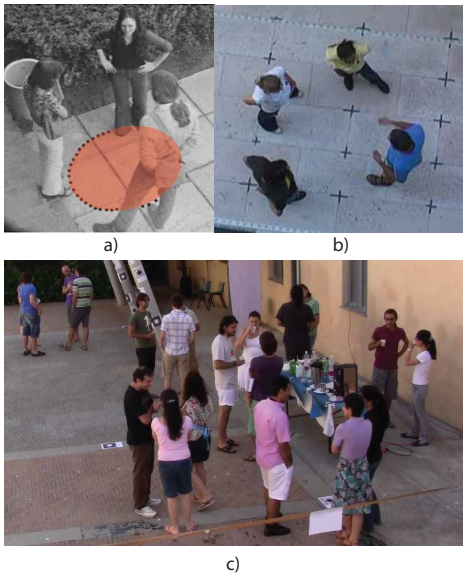

Figure 1: **Examples of F-formations.** a) in orange, the o-space; b) an aerial image of a circular F-formation; c) a party, something similar to a typical surveillance setting with the camera located 2-3 meters from the floor: detecting F-formations here is challenging.

Detecting free-standing conversational groups is useful in many contexts. In video-surveillance, automatically understanding the network of social relationships observed in an ecological scenario may result beneficial for advanced suspect profiling, improving and automatizing SPOT (Screening Passengers by Observation Technique) protocols [18], which nowadays are performed uniquely by human operators.

A robust FCG detector may also impact the social robotics field, where the approaches so far implemented work on few number of people, usually focusing on a single F-formation [32, 65, 51].

Efficient identification of FCGs could be of use in multimedia applications, and especially in semantic tagging [20, 46], where groups of people are currently inferred by the proximity of their faces in the image plane. Adopting systems for 3D pose estimation from 2D images [3] plus an FCG detector could in principle lead to more robust estimations. In this scenario, the extraction of social relationships could help in inferring personality traits [54, 27] and triggering friendship invitation mechanisms [40].

In computer-supported cooperative work (CSCW), being capable of automatically detecting FCG could be a step ahead in understanding how computer systems can support socialization and collaborative activities: e.g., [60, 50, 47, 2]; in this case, FCGs are usually found by hand, or employing wearable sensors.

Manual detection of FCGs occurs also in human computer interaction, for the design of devices reacting to a situational change [29, 55]: here the benefit of the automation of the detection process may lead to a genuine systematic study of how proxemic factors shape the usability of the device.

The last three years have seen works that automatically detect F-formations: Bazzani et al. [6] first proposed the use of positional and orientational information to capture Steady Conversational Groups (SCG); Cristani et al. [16] designed a sampling technique to seek F-formations centres by performing a greedy maximization in a Hough voting space; Hung and Kröse [31] detected F-formations by finding distinct maximal cliques in weighted graphs via graph-theoretic clustering; both the techniques were compared by Setti et al. [57]. A multi-scale extension of the Hough-based approach [16] was proposed by Setti et al. [58]. This improved on previous works, by explicitly modeling F-formations of different cardinalities. Tran et al. [62] followed the graph based approach of [31], extending it to deal with video-sequences and recognizing five kinds of activities.

Our proposed approach detects an arbitrary number of F-formations on single images using a monocular camera, by considering as input the position of people on the ground floor, and their orientation, captured as the head and/or body pose. The approach is iterative, and starts by assuming an arbitrarily high number of F-formations: after that, a hill-climbing optimisation alternates between assigning individuals to F-formations using the efficient graph-cut based optimisation [41], and updating the centres of the F-formations, pruning unsupported groups in accordance with a Minimum Description Length prior. The iterations continue until convergence, which is guaranteed.

As a second contribution, we present a novel set of metrics for group detection. This is not constrained to apply to FCG, but to any set of people considered as a whole, thus embracing generic group or crowd tracking scenarios [62].

The fundamental idea is the concept of *tolerance threshold*, which basically regulates the tolerance on individuating groups, allowing some individual com-

ponents to be missed or external people to be added in a group. Thanks to the tolerance threshold, the concepts of *tolerant match*, *tolerant accuracy* and of *precision* and *recall* can be easily derived. Such measures take inspiration from the group match definition, firstly published in a previous work [16] and adopted in many recent group detection [62, 58] and group tracking methods [7] so far: in practice, it corresponds to fix the tolerance threshold to a predefined value.

In this article, we show that, by letting the tolerance threshold change in a continuous way from maximum to minimum tolerance, it is possible to get an informative and compact measure (in the form of area under the curve) that summarises the behaviour of a given detection methodology. In addition, the tolerant match can be applied specifically to groups of a given cardinality, allowing to obtain specific values of accuracy, precision and recall; this highlights the performance of a given approach in a specific scenario, that is, the ability of capturing small or large groups of people.

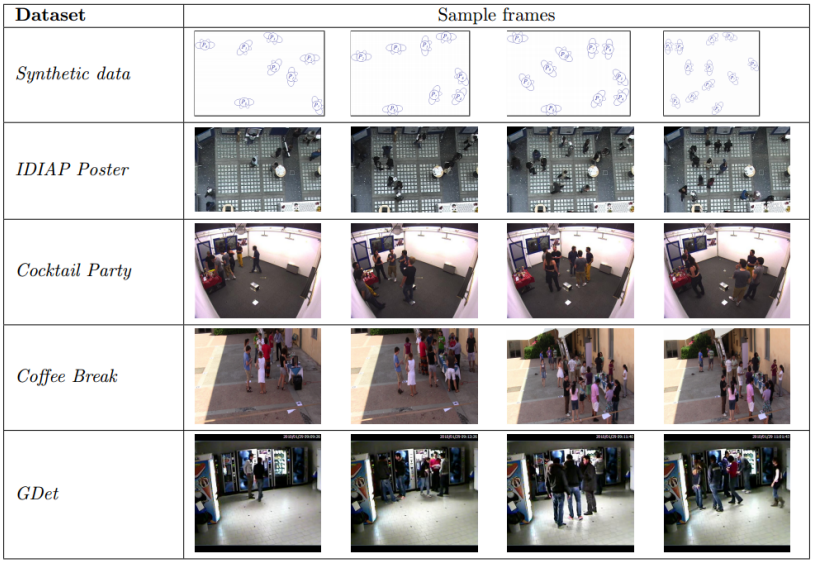

Figure 2: **Sample images of the four real-world datasets.** For each dataset four frames are reported showing different situations of crowd and arrangement.

In the experiments, we apply GCFF to all publicly available datasets (see Fig. 2), consisting of more than 2000 different F-formations over 1024 frames. Comparing against the five most accurate methods in the literature we definitely set the best score on every dataset. In addition, using our novel metrics, we show that GCFF has the best behaviour in terms of robustness to noise, and it is able to capture groups of different cardinalities without changing any threshold. Summarising, the main contributions of this article are the following:

- A novel methodology to detect F-formations from single images acquired by a monocular camera, which operates on positional and orientational information of the individuals in the scene. Unlike previous approaches, our novel methodology is a direct formulation of the sociological principles (proximity, orientation and ease of access) concerning o-spaces. The strong conceptual simplicity and clarity of our approach is an asset in two important ways: we do not require bespoke optimisation techniques, and we make use of established methods known to work reliably and efficiently. Second, and by far more important, the high accuracy and clarity of our approach, along with its basis in sociological principles makes it well suited for use in the social sciences as means of automatically annotating data.

- A rigorous taxonomy of the group entity, which takes from social science and illustrates all the different group manifestations, delineating their main characteristics, in order to go beyond the generic term of group, often misused in the computer vision community.

- A novel set of metrics for group detection, that for the first time models the fact that a group could be partially captured, with some people missing or erroneously taken into account, through the concept of tolerant match. The metrics can be employed to whatever approach involving groups (group tracking included).

The remainder of the paper is organised as follows: the next section presents a literature review of group modeling, with particular emphasis on the terminology adopted, which will be imported from the social and cognitive sciences; the proposed GCFF approach, together with its sociological grounding, is presented afterwards, followed by an extensive experimental evaluation. Finally, we will draw the conclusion and envisage the future perspectives.

—

2 Literature Review

Research on group modeling in computer science is highly multidisciplinary, necessarily encompassing the social and the cognitive sciences when it comes to analyse human interaction. In this multifaceted scenario, characterising the works most related to our approach requires us to distinguish between related sociological concepts; starting with the Goffmanian [25] notions, of (a) "group" vs. "gathering", (b) "social occasion" vs. "social situation", (c) "unfocused" vs. "focused" interaction, and (d) Kendon's [37] specification concerning "common focused" vs. "jointly focused" encounters.

As mentioned in the introduction, *groups* entail some durable membership and organisation, *gatherings* consist of any set of two or more individuals in mutual immediate presence at a given moment. When people are co-present, they tend to behave like one who participates in a social occasion, and the latter provides the structural social context, the general "scheme" or "frame" of

behaviour –like a party, a conference dinner, a picnic, an evening at the theatre, a night in the club, an afternoon at the stadium, a walk together, a day at the office, etc.– within which gatherings (may) develop, dissolve and redevelop in diverse and always different situational social contexts (or social situations, that is, e.g., that specific party, dinner, picnic, etc.) [26].

Unfocused interaction occurs whenever individuals find themselves by circumstance in the immediate presence of others. For instance, when forming a queue or crossing the street at a traffic light junction. On such occasions, simply by virtue of the reciprocal presence, some form of interpersonal communication must take place regardless of individual intent. Conversely, *focused interaction* occurs whenever two or more individuals willingly agree –although such an agreement is rarely verbalised– to sustain for a time a single focus of cognitive and visual attention [25]. Focused gatherings can be further distinguished in *common focused* and *jointly focused* one [37]. The latter entails the sense of a mutual, instead of merely common, activity; a preferential openness to interpersonal communication, an openness one does not necessarily find among strangers at the theatre, for instance; in other words, a special communication license, like in a conversation, a board game, or a joint task carried on by a group of face-to-face interacting collaborators. Participation, in other words, is not at all peripheral but engaged; people are – and display to be – mutually involved [26]. All this can exclude from the gathering others who are present in the situation, as in any FCG at a coffee break with respect to the other ones.

Finally, we should consider the static/dynamic axis concerning the degree of freedom and flexibility of the spatial, positional, and orientational organisation of gatherings. Sometimes, indeed, people maintain approximately their positions for an extended period of time within fixed physical boundaries (e.g., during a meeting); sometimes they move within a delimited area (e.g., at a party); and sometimes they do within a more or less unconstrained space (for instance, people conversing while walking in the street). It is about a continuum, in which we can analytically identify thresholds. Tab. 1 lists some categorised examples of gatherings, considering the taxonomy axis "static/dynamic organisation" and the "unfocused/common-focused/jointly-focused interaction" one. Fig. 3 shows some categorised examples of encounters.

	Unfocused	**Common focused**	**Jointly focused**
Static	open-space offices	conferences, classrooms	meetings, board-game play
Dynamic	waiting rooms, queues, promenades, airports, stations, street-crossings	theatre stands, stadium stands, parades, processions, demonstrations	restaurants, having meal together, work-breaks, bar, clubs, pubs

Table 1: Gatherings categorisation on the basis of focus of attention and spatio-proxemic freedom exemplified by typical social settings/situations.

Within this taxonomy, our interest is on *gatherings*, formed by people *jointly focused* on interacting in a *quasi-static* fashion within a dynamic scenario. Kendon dubbed this scenario as characterising *free-standing conversational groups*, highlighting their spontaneous aggregation/disgregation nature, implying that

	Unfocused	Common focused	Jointly focused
Static			
Dynamic			

Figure 3: **Examples of gatherings categorised by focus of attention and spatio-proxemic freedom.** Jointly focused, dynamic: our case, FCGs at a cocktail party; common focused, dynamic: a parading platoon; unfocused, dynamic: a queue at the airport; jointly focused, static: a meeting; common focused, static: people in a theatre stand; unfocused, static: persons in a waiting room.

their members are jointly focused, and specifying their mainly-static proxemic layout within a dynamic proxemic context.

The following review centres on the case of FCGs and their formation, while for the other cases we refer: with respect to computer vision, to [1] for generic human activity analysis, including single individuals, groups and crowds, and to [33] for a specific survey on crowds; with respect to the sociological literature, to [26] as for unfocused gatherings, to [37, 56] as for common focused ones, and to [36, 49] as for crowds in particular.

The analysis of focused gatherings in computer science had the first traces appearing in the field of human computer interaction and robotics, especially for what concerns context- aware computing, computer-supported cooperative work and social robotics [29, 55, 5, 35]. This happened since the detection of focused gatherings requires finer feature analysis, and in particular body posture inference other than positional cues extraction: these are difficult tasks for traditional computer vision scenarios, where people is captured at low resolution, under diverse illumination conditions, often partially or completely occluded.

In human-computer interaction, F-formation analysis encompasses context-aware computing, by considering spatial relationships among people where space factors become crucial into the design of applications for devices reacting to a situational change [29, 55]. In particular, Ballendat et al. [5] studied how proxemic interaction is expressive when considering cues like position, identity, movement, and orientation. They found that these cues can mediate the simultaneous interaction of multiple people as an F-formation, interpreting and exploiting people's directed attention to other people. So far, the challenge with these applications for researchers has been the hardware design, while the social dynamics are typically not explored. As notable exception, Jungman et

al. [35] studied how different kinds of F-formations (L-shaped vs. face-to-face) identify different kinds of interaction: in particular, they examined whether or not Kendon's observation according to which face-to-face configurations are preferred for competitive interactions whereas L-shaped configurations are associated with cooperative interactions holds in gaming situations. The results partially supported the thesis.

In computer-supported cooperative work, Suzuki and Kato [60] described how different phases of collaborative working were locally and tacitly initiated, accomplished and closed by children by moving back and forth between standing face-to-face formations and sitting screen-facing formations. Morrison et al. [50] studied the impact of the adoption of electronic patient records on the structure of F-formations during hospital ward rounds. Marshall et al. [47] analysed through F-formations the social interactions between visitors and staff in a tourist information centre, describing how the physical structures in the space encouraged and discouraged particular kinds of interactions, and discussing how F-formations might be used to think about augmenting physical spaces. Finally, Akpan et al. [2], for the first time, explored the influence of both physical space and social context (or place) on the way people engage through F-formations with a public interactive display. The main finding is that social properties are more important than merely spatial ones: a conducive social context could overcome a poor physical space and encourage grouping for interaction; conversely, an inappropriate social context could inhibit interaction in spaces that might normally facilitate mutual involvement. So far, no automatic F-formation detection has been applied: positional and orientational information were analysed by hand, while our method is fully automated.

In social robotics, Nieuwenhuisen and Behneke presented Robotinho [51], a robotic tour guide which resembles behaviour of human tour guides and leads people to exhibits in a museum, asking people to come closer and arrange themselves in an F-formation, such that it can attend the visitors adequately while explaining an exhibit. Robotinho detects people by first detecting their faces, and using laser-range measurements to detects legs and trunks. Given this, it is not clear how proper F-formations are recognised. Robotinho essentially improves what has been done by Yousuf et al. [65], that develop a robot that simply detect when an F-formation is satisfied before explaining an exhibit. In this case, F-formations were detected automatically, using advanced sensors (range cameras, etc.) with the possibility of checking just one formation. In our case, a single monocular camera is adopted and the number of F-formations is not bounded.

In computer vision, Groh et al. [28] proposed to use the relative shoulder orientations and distances (using wearable sensors) between each pair of people as a feature vector for training a binary classifier, learning the pairwise configurations of people in a FCG and not. Strangely, the authors discouraged large FCG during the data acquisition, introducing a bias on their cardinality. With our proposal, no markers or positional devices have been considered, and entire FCGs of arbitrary cardinality are found (not pairwise associations only). In his previous work [6], one of the authors started to analyse F-formations by check-

ing the intersection of the view-frustum of neighbouring people, where the view frustum was automatically detected by inferring the head orientation of each single individual in the scene. Under a sociological perspective, the head orientation cue can be exploited as an approximation of a person's focus of visual and cognitive attention, which in turn acts as an indication of the body orientation and the foot position, the last one considered as the most proper way to detect F-formations. Hung and Kröse [31] proposed to consider an F-formation as a dominant-set cluster [52] of an edge-weighted graph, where each node in the graph is a person, and the edges between them measure the affinity between pairs. Such maximal cliques has been defined by Pavan and Pelillo as dominant sets [52], for which a game theoretic approach has been designed to solve the clustering problem under these constraints. More recently, Tran et al. [62] applied a similar graph-based approach for finding groups, which were subsequently analysed by a specific descriptor that encodes people's mutual poses and their movements within the group gathering for activity recognition. In all these three approaches, the common underlying idea is to find set of pairs of individuals with similar mutual pose and orientation, thus considering pairwise proxemics relations as basic elements. This is weak, since in practice it tends to find circular formations (that is, cliques with compact structures), while FCGs have other common layouts (side-by- side, L-shape, etc.). In our case, all kinds of F-formations can be found. In addition, the definition of F-formation requires that no obstacles must invade the o-space (the convex space surrounded by the group members, see Fig. 1a): whereas in the above-mentioned approaches such a condition is not explicitly taken into account, it is a key element in GCFF.

In this sense, GCFF shares more similarities with the work of Cristani et al. [16], where F-formations were found by considering as atomic entity the state of a single person: each individual projects a set of samples in the floor space, that vote for different o-space centres, depending on his or her position and orientation. Votes are then accumulated in a proper Hough space, where a greedy minimization finds the subset of people voting for the same o-space centre, which in turns is free of obstacles. Setti et al. [57] compared the Hough-based approach with the graph-based strategy of Hung and Kröse [31], finding that the former performs better, especially when in presence of high noise. The study was also aimed at analysing how important positional and orientational information are: it turned out that, when in presence of positional information only, the performances of the Hough-based approach decrease strongly, while graph-based approaches are more robust. Another voting-based approach resembling the Hough-based strategy has been designed by Gan et al. [21], who individuated a global interaction space as the overlap area of different individual interaction spaces, that is, conic areas aligned coherently with the body orientations of the interactants (detected using a kinect device). Subsequently, the Hough-based approach has been extended for dealing with groups of diverse cardinalities by Setti et al. [58], who adopted a multi-scale Hough-space, and set the best performance so far.

3 Method

Our approach is strongly based on the formal definition of F-formation given by Kendon [38] (*page 209*):

> *An F-formation arises whenever two or more people sustain a spatial and orientational relationship in which the space between them is one to which they have equal, direct, and exclusive access.*

In particular, an F-formation is the proper organisation of three social spaces: *o-space*, *p-space* and *r-space* (see Fig. 4a).

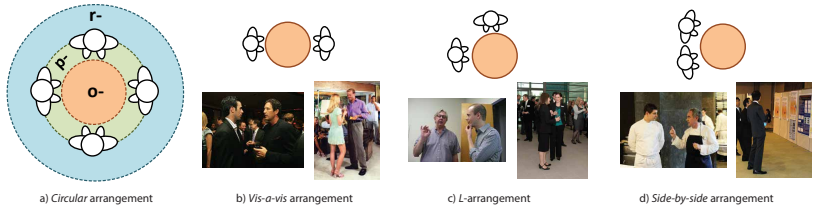

Figure 4: **Structure of an F-formation and examples of F-formation arrangements.** a) Schematization of the three spaces of an F-formation: starting from the centre, *o-space*, *p-space* and *r-space*. b-d) Three examples of F-formation arrangements: for each one of them, one picture highlights the head and shoulder pose, the other shows the lower body posture. For a picture of circular F-formation, see also Fig. 1.

The o-space is a convex empty space surrounded by the people involved in a social interaction, where every participant is oriented inward into it, and no external people are allowed to lie. More in the detail, the o-space is determined by the overlap of those regions dubbed *transactional segments*, where as transactional segment we refer to the area in front of the body that can be reached easily, and where hearing and sight are most effective [13]. In practice, in a F-formation, the transactional segment of a person coincides with the o-space, and this fact has been exploited in our algorithm. The p-space is the belt of space enveloping the o-space, where only the bodies of the F-formation participants (as well as some of their belongings) are placed. People in the p-space participate to an F-formation using the o-space to transmit their messages. The r-space is the space enveloping o- and p-spaces, and is also monitored by the F-formation participants. People joining or leaving a given F-formation mark their arrival as well as their departure by engaging in special behaviours displayed in a special order in special portions of r-space, depending on several factors (context, culture, personality among the others); therefore, here we prefer to avoid the analysis of such complex dynamics, leaving their computational analysis as future work.

F-formations can be organised in different *arrangements*, that is, spatial and orientational layouts (see Fig. 4a-d) [15, 14, 38]. In F-formations of two individuals, usually we have a *vis-a-vis* arrangement, in which the two participants stand and face one another directly; another situation is the *L-arrangement*, when two people lie in a right angle to each other. As studied by Kendon [38], vis-a-vis configurations are preferred for competitive interactions, whereas L-shaped configurations are associated with cooperative interactions. In a *side-by-side* arrangement, people stand close together, both facing the same way; this situation occurs frequently when people stand at the edges of a setting against walls. *Circular* arrangements, finally, hold when F-formations are composed by more than two people; other than being circular, they can assume an approximately linear, semicircular, or rectangular shape.

GCFF finds the o-space of an F-formation, assigning to it those individuals whose transactional segments do overlap, without focusing on a particular arrangement. Given the position of an individual, to identify the transactional segment we exploit orientational information, which may come from the head orientation, the shoulder orientation or the feet layout, in increasing order of reliability [38]. The idea is that the feet layout of a subject indicates the mean direction along which his messages should be delivered, while he is still free to rotate his head and to some extent his shoulders through a considerable arc, before he must begin to turn his lower body as well. The problem is that feet are almost impossible to detect in an automatic fashion, due to the frequent (auto) occlusions; shoulder orientation is also complicated, since most of the approaches of body pose estimation work on 2D data and do not manage auto-occlusions. However, since any sustained head orientation in a given direction is usually associated with a reorientation of the lower body (so that the direction of the transactional segment again coincides with the direction in which the face is oriented [38]), head orientation should be considered proper for detecting transactional segments and, as a consequence, the o-space of an F-formation. In this work, we assume to have as input both positional information and head orientation; this assumption is reasonable due to the massive presence of robust tracking technologies [8] and head orientation algorithms [59, 4, 12].

In addition to this, we consider soft exclusion constraints: in an o-space, F-formation participants should have *equal, direct and exclusive access*. In other words, if person i stands between another person j, and an o-space centre O_g of the F-formation g, this should prevent j from focusing on the o-space, and, as a consequence, from being part of the related F-formation.

In what follows, we formally define the objective function accounting for positional, orientational and exclusion constraints aspects, and show how it can be optimised. Fig. 5 gives a graphical idea of the problem formulation.

3.1 Objective Function

We use $P_i = [x_i, y_i, \theta_i]$ to represent the position x_i, y_i and head orientation θ_i of the individual $i \in \{1, \ldots, n\}$ in the scene. Let TS_i be the a priori distribution which models the transactional segment of individual i. As we explained in the

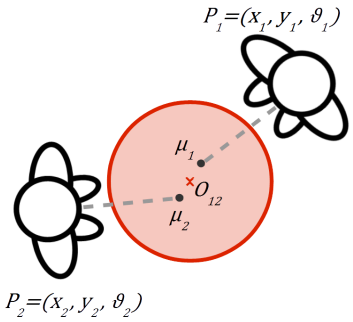

Figure 5: **Schematic representation of the problem formulation.** Two individuals facing each other, the gray dot representing the transitional segment centre, the red cross being the o-space centre and the red area the o-space of the F-formation.

previous section, this segment is coherent with the position and orientation of the head, so we can assume $TS_i \sim \mathcal{N}(\mu_i, \mathbf{\Sigma}_i)$, where $\mu_i = [x_{\mu_i}, y_{\mu_i}] = [x_i + D\cos\theta_i, y_i + D\sin\theta_i]$, $\mathbf{\Sigma}_i = \sigma \cdot \mathbf{I}$ with $\mathbf{I}$ the 2D identity matrix, and D is the distance between the individual i and the centre of its transactional segment (hereafter called *stride*). The stride parameter D can be learned by cross-validation, or fixed a priori accounting for social facts. In practice, we assume the transactional segment of a person having a circular shape, which can be thought as superimposed to the o-space of the F-formation she may be part of.

$O_g = [u_g, v_g]$ indicates the position of a candidate o-space centre for F-formation $g \in \{1, M\}$, while we use G_i to refer to the F-formation containing individual i, considering the F-formation assignment $G_i = g$ for some g. The assignment assumes that each individual i may belong to a single F-formation g only[1] at any given time, and this is reasonable when we are focusing one a single time, that is, an image. It follows naturally the definition of $O_{G_i} = [u_{G_i}, v_{G_i}]$, which represents the position of a candidate o-space centre for an unknown F-formation $G_i = g$ containing i.

At this point, we define the likelihood probability of an individual i's transitional segment centre $C_i = [u_i, v_i]$ given the a priori variable TS_i.

$$\Pr(C_i|TS_i) \propto \exp\left(-\frac{||C_i - \mu_i||_2^2}{\sigma^2}\right) \tag{1}$$

$$= \exp\left(-\frac{(u_i - x_{\mu_i})^2 + (v_i - y_{\mu_i})^2}{\sigma^2}\right) \tag{2}$$

Hence, the probability that an individual i shares an o-space centre O_{G_i} is

[1] For the sake of mathematical simplicity, we assume that each lone individual not belonging to a gathering can be considered as a *spurious* F-formation.

given by

$$\Pr(C_i = O_{G_i}|TS_i) \propto \exp\left(-\frac{(u_{G_i} - x_{\mu_i})^2 + (v_{G_i} - y_{\mu_i})^2}{\sigma^2}\right) \tag{3}$$

and the posterior probability of any overall assignment is given by

$$\Pr(C = O_G|TS) \propto \prod_{i\in[1,n]} \exp\left(-\frac{(u_{G_i} - x_{\mu_i})^2 + (v_{G_i} - y_{\mu_i})^2}{\sigma^2}\right) \tag{4}$$

with C the random variable which models a possible joint location of all the o-space centres, O_G is one instance of this joint location, and TS is the position of all the transitional segments of the individuals in the scene.

Clearly, if the number of o-space centres is unconstrained, the maximum a posteriori probability (MAP) occurs when each individual has his own separate o-space centre, generating a *spurious* F-formation formed by a single individual, that is, $O_{G_i} = TS_i$. To prevent this from happening, we associate a minimum description length prior (MDL) over the number of o-space centres used. This prior takes the same form as dictated by the Akaike Information Criterion (AIC) [10], linearly penalising the log-likelihood for the number of models used.

$$\Pr(C = O_G|TS) \propto \prod_{i\in[1,n]} \exp\left(-\frac{(u_{G_i} - x_{\mu_i})^2 + (v_{G_i} - y_{\mu_i})^2}{\sigma^2}\right) \cdot \exp(-|O_G|) \tag{5}$$

where $|O_G|$ is the number of distinct F-formations.

To find the MAP solution, we take the negative log-likelihood and discarding normalising constants, we have the following objective $J(\cdot)$ in standard form:

$$J(O_G|TS) = \sum_{i\in[1,n]} (u_{G_i} - x_{\mu_i})^2 + (v_{G_i} - y_{\mu_i})^2 + \sigma^{-2}|O_G| \tag{6}$$

As such, this can be seen as optimising a least-squares error combined with an MDL prior. In principle this could be optimised using a standard technique such as k-means clustering combined with a brute force search over all possible choices of k to optimise the MDL cost. In practice, k-means frequently gets stuck in local optima[2] and instead we make use of the graph-cut based optimisation described in [41], and widely used in computer vision [9, 45, 11, 64]

In short, we start from an abundance of possible o-space centres, and then we use a hill-climbing optimisation that alternates between assigning individuals to o-space centres using the efficient graph-cut based optimisation [41] that directly minimises the cost (6), and then minimising the least squares component by updating o-space centres to the mean of O_g, for all the individuals $\{i\}$ currently assigned to the F-formation. The whole process is iterated until convergence.

[2] In fact, using the technique described the least squares component of the error frequently increases, instead of decreasing, as k increases.

This approach is similar to the standard k-means algorithm, sharing both the assignment, and averaging step. However, as the graph-cut algorithm selects the number of clusters, we can avoid local minima by initialising with an excess of model proposals. In practice, we start from the previously mentioned trivial solution in which each individual is associated with its own o-space centre, centred on his position.

Algorithm 1 Finding shared focal centres

```
Initialise with O_{G_i} = TS_i    ∀i ∈ [1, ..., n]
old_cost= ∞
while J(O_G, TS) <old_cost do
  old_cost ← J(O_G, TS)
  run graph cuts to minimise cost (6)
  for ∀g ∈ [1, ..., M] do
    if g is not empty then
      update O_G ← (Σ_{i∈g} TS_i) / |g|
    end if
  end for
end while
```

3.2 Visibility constraints

Finally, we add the natural constraint that people can only join an F-Formation if they can see the o-space centres. By allowing other people to occlude the o-space centre, we are able to capture more subtle nuances such as people being crowded out of F-formations or deliberately ostracised. Broadly speaking, an individual is excluded from an F-formation when another individual stands between him and the group centre. Taking $\theta^g_{i,j}$ as the angle between two individuals about a given o-space centre O_g for which is assumed $G_i = G_j = g$ and d^g_i, d^g_j as the distance of i, or j, respectively from the o-space centre O_g, the following cost captures this property:

$$R_{i,j}(g) = \begin{cases} 0 & \text{if } \theta^g_{i,j} \leq \hat{\theta}, \text{ or } d^g_i < d^g_j \\ \exp\left(K\cos(\theta^g_{i,j})\right)\frac{d^g_i - d^g_j}{d^g_j} & \text{otherwise.} \end{cases} \tag{7}$$

and use the new cost function:

$$J'(O_G|TS) = J(O_G|TS) + \sum_{i,j\in P} R_{i,j}(G_i) \tag{8}$$

$R_{i,j}(g_i)$ acts as a visibility constraint on i regardless of the group person j is assigned to, as such it can be treated as a unary cost or data-term and included in the graph-cut based part of the optimisation. Now we turn to other half of the optimisation - updating the o-space centres. Although, given an assignment of people to a o-space centre, a local minima can be found using any off the shelf

non-convex optimisation, we take a different approach. There are two points to be aware of: first, the difference between J' and J is sharply peaked and close to zero in most locations, and can generally be safely ignored; second and more importantly, we may often want to move out of a local minima. If updating an o-space centre results in a very high repulsion cost to one individual, this can often be dealt with by assigning the individual to a new group, and this will result in a lower overall cost, and more accurate labelling. As such, when optimising the o-space centres, we pass two proposals for each currently active model to graph-cuts – the previous proposal generated, and a new proposal based on the current mean of the F-formation. As the graph-cut based optimisation starts from the previous solution, and only moves to lower cost labellings, the cost always decreases and the procedure is guaranteed to converge to a local optimum.

4 Experiments

The experiments section contains the most exhaustive analysis of the group detection methods in still images carried so far in the computer vision literature, to the best of our knowledge.

In the preliminary part, we describe the five publicly available datasets employed as benchmark, the six methods taken into account as comparison and the metrics adopted to evaluate the detection performances. Subsequently, we start with an explicative example of how our approach GCFF does work, considering a synthetic scenario taken from the Synthetic dataset. The experiments continue with a comparative evaluation of GCFF on all the benchmarks against all the comparative methods, looking for the best performance of each approach. Here, GCFF definitely outperforms all the competitors, setting in all the cases new state-of-the-art scores. The ability of detecting groups of a given cardinality and a noise robustness analysis conclude the section, further promoting our technique.

4.1 Datasets

Five publicly available datasets are used for the experiments: two from [16] (*Synthetic* and *Coffee Break*), one from [31] (*IDIAP Poster Data*), one from [58] (*Cocktail Party*), and one from [6] (*GDet*). A summary of the dataset features is in Table 2, while a detailed presentation of each dataset follows. All these datasets are publicly available and the participants to the original experiments gave their permission to share the images and video for scientific purposes. In Fig. 2, some frames of all the datasets are shown.

Dataset	Data Type	Detection	Detection Quality
Synthetic	synthetic	–	perfect
IDIAP Poster	real	manual	very high
Cocktail Party	real	automatic	high
Coffee Break	real	automatic	low
GDet	real	automatic	very low

Table 2: Summary of the features of the datasets used for experiments.

Synthetic Data[3]

A psychologist generated a set of 10 diverse situations, each one repeated with minor variations for 10 times, resulting in 100 frames representing different social situations, with the aim to span as many configurations as possible for F-formations. An average of 9 individuals and 3 groups are present in the scene, while there are also individuals not belonging to any group. Proxemic information is noiseless in the sense that there is no clutter in the position and orientation state of each individual.

IDIAP Poster Data (IPD)[4]

Over 3 hours of aerial videos (resolution 654×439px) have been recorded during a poster session of a scientific meeting. Over 50 people are walking through the scene, forming several groups over time. A total of 82 images were selected with the idea to maximise the crowdedness and variance of the scenes. Images are unrelated to each other in the sense that there are no consecutive frames, and the time lag between them prevents to exploit temporal smoothness. As for the data annotation, a total of 24 annotators were grouped into 3-person subgroups and they were asked to identify F-formations and their associates from static images. Each person's position and body orientation was manually labelled and recorded as pixel values in the image plane – one pixel represented approximately 1.5cm. The difficulty of this dataset lies in the fact that a great variety of F-formation typologies are present in the scenario (other than circular, L-shapes, side-by-side are present).

Cocktail Party (CP)[5]

This dataset contains about 30 minutes of video recordings of a cocktail party in a $30m^2$ lab environment involving 7 subjects. The party was recorded using four synchronised angled-view cameras (15Hz, 1024×768px, jpeg) installed in the corners of the room. Subject's positions were logged using a particle filter-based body tracker [42] while head pose estimation is computed as in [43]. Groups in one frame every 5 seconds were manually annotated by an expert, resulting in a total of 320 labelled frames for evaluation. This is the

[3] http://profs.sci.univr.it/{~}cristanm/datasets.html
[4] http://www.idiap.ch/scientific-research/resources
[5] The dataset is available by writing to the authors of [43]

first dataset where proxemic information is estimated automatically, so errors may be present. Anyway, due to the highly supervised scenario, errors are very few.

Coffee Break (CB)[6]

The dataset focuses on a coffee-break scenario of a social event, with a maximum of 14 individuals organised in groups of 2 or 3 people each. Images are taken from a single camera with resolution of 1440×1080px. People positions have been estimated by exploiting multi-object tracking on the heads, and head detection has been performed afterwards [61], considering solely 4 possible orientations (front, back, left and right) in the image plane. The tracked positions and head orientations were then projected onto the ground plane. Considering the ground truth data, a psychologist annotated the videos indicating the groups present in the scenes, for a total of 119 frames split in two sequences. The annotations were generated by analysing each frame in combination with questionnaires that the subjects filled in. This dataset represent one of the most difficult benchmark, since the rough head orientation information, also affected by noise, gives in many cases unreliable information. Anyway, it represents also one of the most realistic scenario, since all the proxemic information comes from automatic, off/the/shelf, computer vision tools.

GDet[7]

The dataset is composed by 5 subsequences of images acquired by 2 angled-view low resolution cameras (352×328px) a number of frames spanning from 17 to 132, for a total of 403 annotated frames. The scenario is a vending machines area where people meet and chat while they are having coffee. This is similar to Coffee Break scenario but in this case the scenario is indoor, which makes occlusions in this case many and severe; moreover, people in this scenario knows each other in advance. The videos were acquired with two monocular cameras, located on opposite angles of the room. To ensure the natural behaviour of people involved, they were not aware of the experiment purposes. Ground truth generations follows the same protocol as in Coffee Break; but in this case people tracking has been performed using the particle filter proposed in [42]. Also in this case, head orientation was fixed to 4 angles. This dataset, together with Coffee Break, is the closest to what computer vision can give as input to our a FCG detection technique.

4.2 Alternative methods

As alternative methods, we consider all the suitable approaches proposed in the state of the art. Six methods are taken into account, one exploiting the concept of *view frustum* (IRPM [6]), two approaches based on dominant-sets (DS [31]

[6] http://profs.sci.univr.it/~cristanm/datasets.html
[7] http://www.lorisbazzani.info/code-datasets/multi-camera-dataset

and IGD [62]) and three different version of Hough Voting approaches using linear accumulator [16], entropic accumulator [57] and a multi-scale procedure [58]. It follows a brief overview of the different methods – some of them being explained in the Introduction and in the Literature Review section. Please refer to the specific papers for more details about the algorithms.

Inter-Relation Pattern Matrix (IRPM)

Proposed by Bazzani et al. [6], it uses the head direction to infer the 3D view frustum as approximation of the Focus of Attention (FoA) of an individual; given the FoA and proximity information, interactions are estimated: the idea is that close-by people whose view frustum is intersecting are in some way interacting.

Dominant Sets (DS)

Presented by Hung and Kröse [31], this algorithm considers an F-formation as a dominant-set cluster [52] of an edge-weighted graph, where each node in the graph is a person, and the edges between them measure the affinity between pairs.

Interacting Group Discovery (IGD)

Presented by Tran et al. [62], it is based on dominant sets extraction from an undirected graph where nodes are individuals and the edges have a weight proportional to how much people are interacting. This method is similar to DS, but it differs in the way the weights of the edges in the graph are computed; in particular, it exploits social cues to compute this weight, approximating the attention of an individual as an ellipse centred at a fixed offset in front of him. Interaction is based on the intersection of the attention ellipses related to two individuals: the more overlap between ellipses, the more they are interacting.

Hough Voting for F-formation (HVFF)

Under this caption, we consider a set of methods based on a Hough Voting strategy to build accumulation spaces and find local maxima of this function. The general idea is that each individual is associated with a Gaussian probability density function which describes the position of the o-space centre he is pointing at. The pdf is approximated by a set of samples, which basically vote for a given o-space centre location. The voting space is then quantized and the votes are aggregated on squared cells, so to form a discrete accumulation space. Local maxima in this space identify o-space centres, and consequently, F-formations. The first work in this field is [16], where the votes are linearly accumulated by just summing up all the weights of votes belonging to the same cell. A first improvement of this approach is presented in [57], where the votes are aggregated by using the weighted Boltzmann entropy function. In [58] a multi-scale approach is used on top of the entropic version: the idea is that groups with higher cardinality tends to arrange around a larger o-space; the entropic group

search runs for different o-space dimensions by filtering groups cardinalities; afterwards, a fusion step is based on a majority criterion.

4.3 Evaluation metrics

As accuracy measures, we adopt the metrics proposed in [16] and extended in [57]: we consider a group as correctly estimated if at least $\lceil (T \cdot |G|) \rceil$ of their members are found by the grouping method and correctly detected by the tracker, and if no more than $1 - \lceil (T \cdot |G|) \rceil$ false subjects (of the detected tracks) are identified, where $|G|$ is the cardinality of the labelled group G, and $T \in]0, 1]$ is an arbitrary threshold, called *tolerance threshold*. In particular, we focus on two interesting values of T: $2/3$ and 1.

With this definition of *tolerant match*, we can determine for each frame the correctly detected groups (true positives – TP), the miss-detected groups (false negatives – FN) and the hallucinated groups (false positives – FP). With this, we compute the standard pattern recognition metrics precision and recall:

$$precision = \frac{TP}{TP + FP} \quad , \quad recall = \frac{TP}{TP + FN} \tag{9}$$

and the F_1 score defined as the harmonic mean of precision and recall:

$$F_1 = 2 \cdot \frac{precision \cdot recall}{precision + recall} \tag{10}$$

In addition to these metrics, we present in this paper a new metric which is independent from the tolerance threshold T. We compute this new score as the area under the curve (AUC) in the F_1 vs. T graph with T varying from $1/2$ to 1[8]. We will call it *Global Tolerant Matching* score (GTM). Since in our experiments we only have groups up to 6 individuals, without loss of generality we consider T varying with 3 equal steps in the range stated above.

Moreover, we will discuss results also in terms of group cardinality, by computing the F_1 score for each cardinality separately and then computing mean and standard deviation.

4.4 An explicative example

Figure 6 gives a visual insight of our graph-cuts process. Given the position and orientation of each individual P_i, the algorithm starts by computing the transitional segments C_i. At the first iteration 0, the candidate o-space centres O_i are initialized, and are coincident with the transitional segments C_i; in this example are present 11 individuals, so 11 candidate o-space centres are generated. After iteration 1, the proposed segmentation process provides 1 singleton (P_{11}) and 5 FCGs of two individuals each. We can appreciate different configurations such

[8] Please note that we avoid to consider $0 < T < 1/2$, since in this range we are accepting as good those groups where more than the half of the subjects is missing or false positive, resulting in useless estimates.

as *vis-a-vis* ($O_{1,2}$), L-shape ($O_{3,4}$) and side-by-side ($O_{5,6}$). Still, the grouping in the bottom part of the image is wrong (P_7 to P_{10}), since it violates the exclusion principle. In iteration 2, the previous candidate o-space centres is considered as initialization, and a new graph is built. In this new configuration, the group $O_{7,10}$ is recognized as violating the visibility constraint and thus the related edge is penalized; a new run of graph-cuts minimization allows to correctly cluster the FCGs in a singleton (P_{10}) and a FCG formed by three individuals ($O_{7,8,9}$), which corresponds to the ground truth (visualized as the dashed circles).

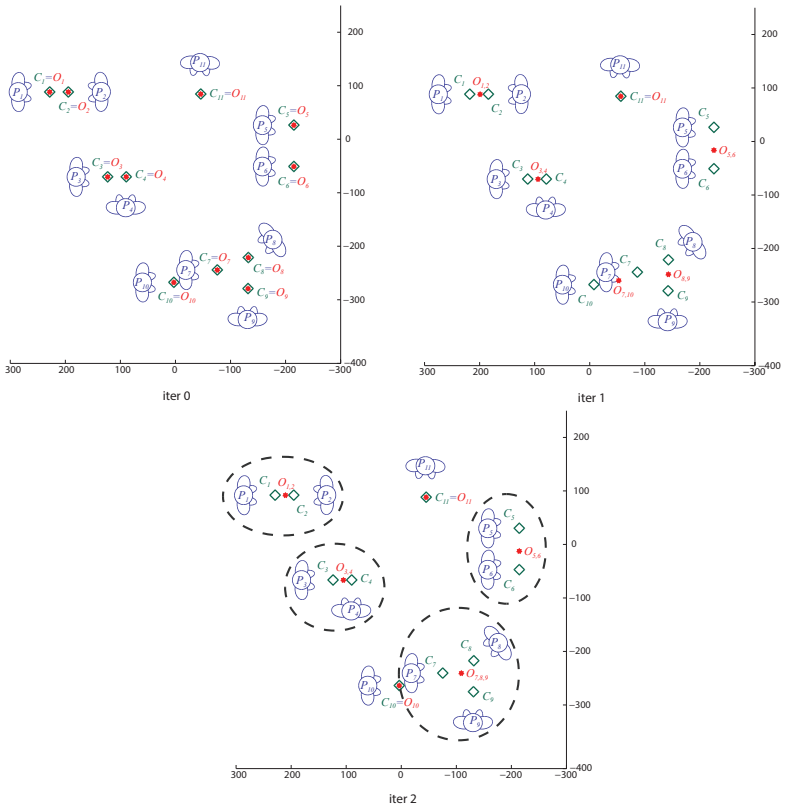

Figure 6: **An explicative example.** Iteration 0: initialization with the candidate o-space centres $\{O\}$ coincident with the transitional segment of each individual $\{C\}$. Iteration 1: first graph-cuts run; easy groups are correctly clustered while the most complex still present errors (the FCG formed by P_7 and P_20 violates the visibility constraint). Iteration 2: the second graph-cuts run correctly detects the $O_{7,8,9}$ F-formation (at the bottom). Se text for more details.

4.5 Best results analysis

Given the metrics explained above, the first test analyses the best performances for each method on each dataset; in practice, a tuning phase has been carried out for each method/dataset combination in order to get the best performances[9]. Best parameters (found on half of one sequence by cross-validation, and kept unchanged for the remaining datasets) are reported in Table 3. Please note, finding the right parameters can also fixed by hand, since the stride D depends on the social context under analysis (formal meetings will have higher D, the presence of tables and similar items may also increase the diameter of the FCGs): with a given D, for example, it is assumed that circular F-formations will have diameter of $2D$. The parameter σ indicates how much we are permissive in accepting deviations from such a diameter. Moreover, D depends also on the different measure units (pixels/cm) which characterize the proxemic information associated to each individual in the scene.

Dataset	stride D	*std* σ
Synthetic	30	80
IDIAP Poster	20	45
Cocktail Party	70	170
Coffee Break	30	85
GDet	30	200

Table 3: Parameters used in the experiments for each dataset. These parameters are the results of a tuning phase and the difference are due to different measure units (pixels/cm) and different social environments (indoor/outdoor, formal/informal, etc.).

Table 4 shows best results by considering the threshold $T = 2/3$, which corresponds to find at least 2/3 of the members of a group, no more than 1/3 of false subjects; while Table 5 presents results with $T = 1$, considering a group as correct if all and only its members are detected. The proposed method outperforms all the competitors, on all the datasets. With $T = 2/3$, three observations can be made: the first is that our approach GCFF improves substantially the precision (of 13% in average) and even more definitely the recall scores (of 17% in average) of the state of the art approaches. The second is that our approach produces the same score for both the precision and the recall; this is very convenient and convincing, since so far all the approaches of FCG detections have shown to be weak in the recall dimension. The third observation is that GCFF performs well both in the case where no errors in the position or orientation of the people are present (as the Synthetic dataset) and in the cases where strong noise of position and orientation is present (Coffee

[9]We did not have code for Dominant Sets [31] and thus we used results provided directly from the authors of the method for a subset of data. For this reason, average results over all the datasets are only averaged over 3 datasets, and cannot be taken into account for a fair comparison.

	Synthetic			*IDIAP Poster*			*Cocktail Party*			*Coffee Break*			*GDet*			***Total***		
	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1
IRPM [6]	0.85	0.80	0.82	0.82	0.74	0.78	0.56	0.43	0.49	0.68	0.50	0.57	0.77	0.47	0.58	0.70	0.49	0.56
DS [31]	0.85	0.97	0.90	0.91	0.92	0.91	–	–	–	0.69	0.65	0.67	–	–	–	*0.81*[9]	*0.83*[9]	*0.82*[9]
IGD [62]	0.95	0.71	0.81	0.80	0.68	0.73	0.81	0.61	0.70	0.81	0.78	0.79	0.83	0.36	0.50	0.68	0.76	0.70
HVFF lin [16]	0.75	0.86	0.80	0.90	0.95	0.92	0.59	0.74	0.65	0.73	0.86	0.79	0.66	0.68	0.67	0.75	0.79	0.76
HVFF ent [57]	0.79	0.86	0.82	0.86	0.89	0.87	0.78	0.83	0.80	0.76	0.86	0.81	0.69	0.71	0.70	0.78	0.78	0.77
HVFF ms [58]	0.90	0.94	0.92	0.87	0.91	0.89	0.81	0.81	0.81	0.83	0.76	0.79	0.71	0.73	0.72	0.84	0.66	0.74
GCFF	**0.97**	**0.98**	**0.97**	**0.94**	**0.96**	**0.95**	**0.84**	**0.86**	**0.85**	**0.85**	**0.91**	**0.88**	**0.92**	**0.88**	**0.90**	**0.89**	**0.89**	**0.89**

Table 4: Average precision, recall and F_1 scores for all the methods and all the datasets. ($T = 2/3$)

	Synthetic			*IDIAP Poster*			*Cocktail Party*			*Coffee Break*			*GDet*			***Total***		
	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1	prec.	rec.	F_1
IRPM [6]	0.53	0.47	0.50	0.71	0.64	0.67	0.28	0.17	0.21	0.27	0.23	0.25	0.59	0.29	0.39	0.46	0.29	0.35
DS [31]	0.68	0.80	0.74	0.79	0.82	0.81	–	–	–	0.40	0.38	0.39	–	–	–	*0.60*[9]	*0.63*[9]	*0.62*[9]
IGD [62]	0.30	0.22	0.25	0.31	0.27	0.29	0.23	0.10	0.13	0.50	0.50	0.50	0.67	0.20	0.31	0.45	0.21	0.27
HVFF lin [16]	0.64	0.73	0.68	0.80	0.86	0.83	0.26	0.27	0.27	0.41	0.47	0.44	0.43	0.45	0.44	0.43	0.46	0.44
HVFF ent [57]	0.47	0.52	0.49	0.72	0.74	0.73	0.28	0.30	0.29	0.47	0.52	0.49	0.44	0.45	0.45	0.42	0.44	0.43
HVFF ms [58]	0.72	0.73	0.73	0.73	0.76	0.74	0.30	0.30	0.30	0.40	0.38	0.39	0.44	0.45	0.45	0.44	0.45	0.45
GCFF	**0.91**	**0.91**	**0.91**	**0.85**	**0.87**	**0.86**	**0.63**	**0.65**	**0.64**	**0.61**	**0.64**	**0.63**	**0.73**	**0.68**	**0.71**	**0.71**	**0.70**	**0.71**

Table 5: Average precision, recall and F_1 scores for all the methods and all the datasets. ($T = 1$)

Break, GDet).

When moving to tolerance threshold equal to 1 (all the people in a group have to be individuated, and no false positive are allowed) the performance is reasonably lower, but the increment is even stronger w.r.t. to the state of the art, in general on all the datasets: in particular, on the Cocktail Party dataset, the results are more than twice the scores of the competitors. Finally, even in this case, GCFF produces a very similar score for precision and recall.

A performance analysis is also provided by changing the tolerance threshold T. Fig. 7 shows the average F_1 scores for each method computed over all the frames and datasets. From the curves we can appreciate how the proposed method is consistently best performing for each T-value. In the legend of Fig. 7 the Global Tolerant Matching score is also reported. Again, GCFF is outperforming the state of the art, independently from the choice of T.

The reason why our approach does better than the competitors has been explained in the state of the art section, here briefly summarized: the Dominant Set-based approaches DS and IGD, even if they are based on an elegant optimization procedure, tend to find circular groups, and are weaker in individuating other kinds of F-formations. Hough-based approaches HVFF X (X= *lin*, *ent*, *ms*) have a good modeling of the F-formation, allowing to find any shape, but rely on a greedy optimization procedure. Finally, IRPM approach has a rough modeling of the F-formation. Our approach viceversa has a rich modeling of the F-formation, and a powerful optimization strategy.

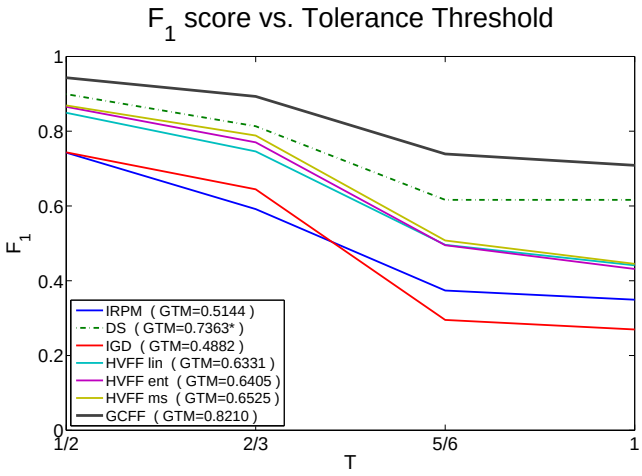

Figure 7: **Global F_1 score vs. tolerance threshold T.** Between brackets in legend the Global Tolerant Matching score. Dominant Sets (DS) is averaged over 3 datasets only, because of results availability. (Best viewed in colour)

4.6 Cardinality analysis

As stated in [58], some methods are shown to work better with some group cardinalities. In this experiment, we sistematically check this aspect, evaluating the performance of all the considered methods in individuating groups with a particular number of individuals. Since Synthetic, Coffee Break and IDIAP Poster Session datasets only have groups of cardinality 2 and 3, we only focus on the remaining 2 datasets, which have a more uniform distribution of groups cardinalities. Tables 6 and 7 show F_1 scores for each method and each group cardinality respectively for Cocktail Party and GDet datasets. In both cases the proposed method outperforms the other state of the art methods in terms of higher average F_1 score, with very low standard deviation. In particular, only IRPM gives in GDet dataset results which are more stable than ours, but they are definitely poorer.

4.7 Noise analysis

In this experiment, we show how the methods behave against different degrees of clutter. For this sake, we consider the Synthetic dataset as starting point and we add to the proxemic state of each individual of each frame some random values based on a known noise distribution. We assume that the noise follows a Gaussian distribution with mean 0, and noise on each dimension (position, orientation) is uncorrelated. For our experiments we used $\sigma_x = \sigma_y = 20$cm and $\sigma_\theta = 0.1$rad. In our experiments, we consider 11 levels of noise $L_n = 0, \ldots, 10$,

	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	**Avg**	**Std**
# groups	81	82	44	55	147	–	–
IRPM [6]	0.26	0.53	0.74	0.42	0.59	0.51	0.18
IGD [62]	0.06	0.52	0.66	0.73	0.85	0.56	0.30
HVFF lin [16]	0.38	*0.76*	0.57	0.67	0.94	0.66	0.21
HVFF ent [57]	0.45	0.75	0.69	0.73	*0.96*	0.71	0.18
HVFF ms [58]	0.49	0.74	0.70	0.71	*0.96*	0.72	0.17
GCFF	*0.59*	0.64	*0.80*	*0.85*	0.94	**0.76**	**0.14**

Table 6: Cocktail Party – F_1 score vs. cardinality. ($T = 1$)

	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	**Avg**	**Std**
# groups	197	124	22	35	13	–	–
IRPM [6]	0.40	0.59	0.45	0.42	0.35	0.44	**0.09**
IGD [62]	0.15	0.52	0.33	0.54	0.83	0.47	0.25
HVFF lin [16]	0.51	0.76	0.03	0.16	0.13	0.32	0.31
HVFF ent [57]	0.57	0.73	0.24	0.23	0.13	0.38	0.26
HVFF ms [58]	0.56	0.78	0.17	0.41	0.67	0.52	0.23
GCFF	*0.74*	*0.87*	*0.53*	*0.77*	*0.88*	**0.76**	0.14

Table 7: GDet– F_1 score vs. cardinality. ($T = 1$)

where

$$\begin{cases} x_n(L_n) = x + \text{randsample}(\mathcal{N}(0, L_n * \sigma_x)) \\ y_n(L_n) = y + \text{randsample}(\mathcal{N}(0, L_n * \sigma_y)) \\ \theta_n(L_n) = \theta + \text{randsample}(\mathcal{N}(0, L_n * \sigma_\theta)) \end{cases} \tag{11}$$

In particular, we produce results by adding noise on position only (leaving the orientation at its exact value), on orientation only (leaving the position of each individual at its exact value) and on both position and orientation. Fig. 8 shows F_1 scores for each method while increasing the noise level. In this case we can appreciate that with high orientation and combined noise IGD performs comparably or better than GCFF; this is a confirmation of the fact that methods based on Dominant Sets are performing very well when the orientation information is not reliable, as already stated in [57].

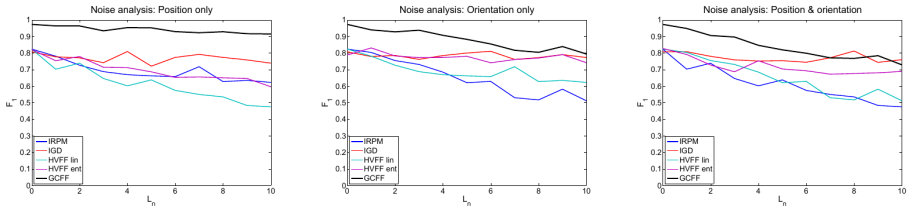

Figure 8: **Noise analysis.** F_1 score vs. Noise Level on position (left), orientation (centre) and combined (right). (Best viewed in colour)

5 Conclusions

In this paper we presented a statistical framework for the detection of free-standing conversational groups (FCG) in still images. FCGs represent very common and crucial social events, where social ties (intimate vs. formal relationships) pop out naturally; for this reason, detection of FCGs is of primary importance in a wide spectra of application. The proposed algorithm is based on a graph-cuts minimization scheme, which essentially clusters individuals into groups; in particular, the computational model implements the sociological definition of F-formation, describing how people forming a FCG will locate in the space. The take-home message is that having basic proxemic information (people location and orientation) is enough to individuate groups with high accuracy. This claim originates from one of the most exhaustive experimental session implemented so far on this matter, with 5 diverse datasets taken into account, and all the best approaches in the literature considered as competitors; in addition to this, a deep analysis on the robustness to noise and on the capability of individuating groups of a given cardinality have been also carried out. The natural extension of this study consists in analyzing the temporal information, that is, video sequences: in this scenario, interesting phenomena such as entering or exiting a group could be considered and modeled, and the temporal smoothness can be exploited to generate even more precise FCG detections.

Acknowledgments

F. Setti and C. Bassetti are supported by the VisCoSo project grant, financed by the Autonomous Province of Trento through the "Team 2011" funding programme. Portions of the research in this paper used the Idiap Poster Data Corpus made available by the Idiap Research Institute, Martigny, Switzerland.

References

[1] Jake K. Aggarwal and Michael S. Ryoo. Human activity analysis: A review. *ACM Computing Surveys (CSUR)*, 43(3):16:1–16:43, 2011. 8

[2] Imeh Akpan, Paul Marshall, Jon Bird, and Daniel Harrison. Exploring the effects of space and place on engagement with an interactive installation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2213–2222, 2013. 4, 9

[3] Mykhaylo Andriluka and Leonid Sigal. Human context: Modeling human-human interactions for monocular 3d pose estimation. In *Articulated Motion and Deformable Objects*, pages 260–272. Springer, 2012. 4

[4] Sileye O. Ba and Jean-Marc Odobez. Multiperson visual focus of attention from head pose and meeting contextual cues. *IEEE Transactions on Pat-*

tern Analysis and Machine Intelligence (PAMI), 33(1):101–116, January 2011. 12

[5] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. Proxemic interaction: designing for a proximity and orientation-aware environment. In *ACM International Conference on Interactive Tabletops and Surfaces*, pages 121–130, 2010. 8

[6] Loris Bazzani, Marco Cristani, Diego Tosato, Michela Farenzena, Giulia Paggetti, Gloria Menegaz, and Vittorio Murino. Social interactions by visual focus of attention in a three-dimensional environment. *Expert Systems*, 30(2):115–127, May 2013. 2, 4, 9, 16, 18, 19, 23, 25

[7] Loris Bazzani, Vittorio Murino, and Marco Cristani. Decentralized particle filter for joint individual-group tracking. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1886–1893, 2012. 2, 5

[8] Ben Benfold and Ian Reid. Stable multi-target tracking in real-time surveillance video. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3457–3464, 2011. 12

[9] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 23(11):1222–1239, November 2001. 14

[10] Hamparsum Bozdogan. Model selection and Akaike's Information Criterion (AIC): The general theory and its analytical extensions. *Psychometrika*, 52(3):345–370, 1987. 14

[11] Neill D. F. Campbell, George Vogiatzis, Carlos Hernández, and Roberto Cipolla. Automatic 3D object segmentation in multiple views using volumetric graph-cuts. *Image and Vision Computing*, 28(1):14–25, 2010. 14

[12] Cheng Chen and Jean-Marc Odobez. We are not contortionists: Coupled adaptive learning for head and body orientation estimation in surveillance video. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1544–1551, 2012. 12

[13] T. Matthew Ciolek. The proxemics lexicon: A first approximation. *Journal of Nonverbal Behavior*, 8(1):55–79, 1983. 11

[14] T. Matthew Ciolek and Adam Kendon. Environment and the spatial arrangement of conversational encounters. *Sociological Inquiry*, 50(3-4):237–271, 1980. 3, 12

[15] Mark Cook. Experiments on orientation and proxemics. *Human Relations*, 23(1):61–76, 1970. 12

[16] Marco Cristani, Loris Bazzani, Giulia Paggetti, Andrea Fossati, Diego Tosato, Alessio Del Bue, Gloria Menegaz, and Vittorio Murino. Social interaction discovery by statistical analysis of f-formations. In *British Machine Vision Conference (BMVC)*, pages 23.1–23.12, 2011. 1, 2, 4, 5, 10, 16, 19, 20, 23, 25

[17] Marco Cristani, Vittorio Murino, and Alessandro Vinciarelli. Socially intelligent surveillance and monitoring: Analysing social dimensions of physical space. In *IEEE International Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 51–58, 2010. 1

[18] Justin Florence and Robert Friedman. Profiles in terror: A legal framework for the behavioral profiling paradigm. *George Mason Law Review*, 17:423–481, 2009. 3

[19] Donelson R. Forsyth. *Group Dynamics*. Cengage Learning, 2010. 2

[20] Andrew C. Gallagher and Tsuhan Chen. Understanding images of groups of people. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 256–263, 2009. 4

[21] Tian Gan, Yongkang Wong, Daqing Zhang, and Mohan S. Kankanhalli. Temporal encoded f-formation system for social interaction detection. In *ACM International Conference on Multimedia*, pages 937–946, 2013. 10

[22] Carolina Gárate, Sofia Zaidenberg, Julien Badie, and François Brémond. Group tracking and behavior recognition in long video surveillance sequences. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, volume 2, pages 396–402, 2014. 2

[23] Daniel Gatica-Perez. Automatic nonverbal analysis of social interaction in small groups: A review. *Image and Vision Computing*, 27(12):1775–1787, 2009. 1, 2

[24] Weina Ge, Robert T. Collins, and R. Barry Ruback. Vision-based analysis of small groups in pedestrian crowds. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 34(5):1003–1016, 2012. 1, 2

[25] Erving Goffman. *Encounters: Two studies in the sociology of interaction*. Bobbs-Merrill, 1961. 2, 6, 7

[26] Erving Goffman. *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Free Press, 1966. 2, 3, 7, 8

[27] Samuel D. Gosling, Sam Gaddis, and Simine Vazire. Personality impressions based on facebook profiles. In *International Conference on Weblogs and Social Media (ICWSM)*, volume 7, pages 1–4, 2007. 4

[28] Georg Groh, Alexander Lehmann, Jonas Reimers, Marc René Frieß, and Loren Schwarz. Detecting social situations from interaction geometry. In *IEEE International Conference on Social Computing (SocialCom)*, pages 1–8, 2010. 9

[29] Eva Hornecker. A design theme for tangible interaction: embodied facilitation. In *European Conference on Computer-Supported Cooperative Work (ECSCW)*, pages 23–43, 2005. 4, 8

[30] Hayley Hung and Daniel Gatica-Perez. Estimating cohesion in small groups using audio-visual nonverbal behavior. *IEEE Transactions on Multimedia*, 12(6):563–575, 2010. 2

[31] Hayley Hung and Ben Kröse. Detecting f-formations as dominant sets. In *International Conference on Multimodal Interfaces (ICMI)*, pages 231–238, 2011. 2, 4, 10, 16, 18, 19, 22, 23

[32] Helge Hüttenrauch, Kerstin Severinson Eklundh, Anders Green, and Elin Anna Topp. Investigating spatial relationships in human-robot interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5052–5059, 2006. 4

[33] Julio Cezar Silveira Jacques Junior, Soraia Raupp Musse, and Claudio Rosito Jung. Crowd analysis using computer vision techniques. *IEEE Signal Processing Magazine*, 27(5):66–77, 2010. 8

[34] Dinesh Babu Jayagopi, Hayley Hung, Chuohao Yeo, and Daniel Gatica-Perez. Modeling dominance in group conversations using nonverbal activity cues. *IEEE Transactions on Audio, Speech, and Language Processing*, 17(3):501–513, 2009. 2

[35] Manuela Jungmann, Richard Cox, and Geraldine Fitzpatrick. Spatial play effects in a tangible game with an f-formation of multiple players. In *Australasian User Interface Conference (AUIC)*, volume 150, pages 57–66, 2014. 8, 9

[36] Naz Kaya and Brigitte Burgess. Territoriality seat preferences in different types of classroom arrangements. *Environment and Behavior*, 39(6):859–876, 2007. 8

[37] Adam Kendon. Goffman's approach to face-to-face interaction. *Erving Goffman: Exploring the interaction order*, pages 14–40, 1988. 6, 7, 8

[38] Adam Kendon. *Conducting interaction: Patterns of behavior in focused encounters*. Cambridge University Press, 1990. 2, 3, 11, 12

[39] Adam Kendon. *The negotiation of context in face-to-face interaction*. Cambridge University Press, 1992. 3

[40] Michal Kosinski, David Stillwell, and Thore Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences (PNAS)*, 110(15):5802–5805, 2013. 4

[41] L'ubor Ladickỳ, Chris Russell, Pushmeet Kohli, and Philip H. S. Torr. Inference methods for crfs with co-occurrence statistics. *International Journal of Computer Vision*, 103(2):213–225, 2013. 4, 14

[42] Oswald Lanz. Approximate bayesian multibody tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 28(9):1436–1449, 2006. 17, 18

[43] Oswald Lanz and Roberto Brunelli. Joint bayesian tracking of head location and pose from low-resolution video. In *Multimodal Technologies for Perception of Humans*, pages 287–296. Springer, 2008. 17

[44] Boris Lau, Kai O. Arras, and Wolfram Burgard. Multi-model hypothesis group tracking and group size estimation. *International Journal of Social Robotics*, 2(1):19–30, 2010. 2

[45] Herve Lombaert, Yiyong Sun, Leo Grady, and Chenyang Xu. A multilevel banded graph cuts method for fast image segmentation. In *International Conference on Computer Vision (ICCV)*, pages 259–265, 2005. 14

[46] Manuel J. Marin-Jimenez, Andrew Zisserman, and Vittorio Ferrari. "here's looking at you, kid." detecting people looking at each other in videos. In *British Machine Vision Conference (BMVC)*, pages 1–12, 2011. 4

[47] Paul Marshall, Yvonne Rogers, and Nadia Pantidi. Using f-formations to analyse spatial patterns of interaction in physical environments. In *ACM Conference on Computer Supported Cooperative Work*, pages 445–454, 2011. 4, 9

[48] Riccardo Mazzon, Fabio Poiesi, and Andrea Cavallaro. Detection and tracking of groups in crowd. In *IEEE International Conference on Advanced Video and Signal-based Surveillance (AVSS)*, pages 202–207, 2013. 2

[49] Clark McPhail. From clusters to arcs and rings. *Research in Community Sociology*, 1:35–57, 1994. 8

[50] Cecily Morrison, Matthew Jones, Alan Blackwell, and Alain Vuylsteke. Electronic patient record use during ward rounds: A qualitative study of interaction between medical staff. *Critical Care*, 12(6):R148, 2008. 4, 9

[51] Matthias Nieuwenhuisen and Sven Behnke. Human-like interaction skills for the mobile communication robot robotinho. *International Journal of Social Robotics*, 5(4):549–561, 2013. 4, 9

[52] Massimiliano Pavan and Marcello Pelillo. Dominant sets and pairwise clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(1):167–172, 2007. 10, 19

[53] Zhen Qin and Christian R. Shelton. Improving multi-target tracking via social grouping. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1972–1978, 2012. 2

[54] John Raacke and Jennifer Bonds-Raacke. Myspace and facebook: Applying the uses and gratifications theory to exploring friend-networking sites. *Cyberpsychology & Behavior*, 11(2):169–174, 2008. 4

[55] Holger Schnädelbach. Hybrid spatial topologies. *The Journal of Space Syntax*, 3(2):204–222, 2012. 4, 8

[56] David Schweingruber and Clark McPhail. A method for systematically observing and recording collective action. *Sociological Methods Research*, 27(4):451–498, 1999. 8

[57] Francesco Setti, Hayley Hung, and Marco Cristani. Group detection in still images by f-formation modeling: A comparative study. In *International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)*, pages 1–4, 2013. 2, 4, 10, 19, 20, 23, 25

[58] Francesco Setti, Oswald Lanz, Roberta Ferrario, Vittorio Murino, and Marco Cristani. Multi-scale f-formation discovery for group detection. In *IEEE International Conference on Image Processing (ICIP)*, 2013. 2, 4, 5, 10, 16, 19, 23, 24, 25

[59] Kevin Smith, Sileye O. Ba, Jean-Marc Odobez, and Daniel Gatica-Perez. Tracking the visual focus of attention for a varying number of wandering people. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 30(7):1212–1229, 2008. 12

[60] Hideyuki Suzuki and Hiroshi Kato. Interaction-level support for collaborative learning: Algoblock – an open programming language. In *International Conference on Computer Support for Collaborative Learning*, pages 349–355, 1995. 4, 9

[61] Diego Tosato, M. Spera, Marco Cristani, and Vittorio Murino. Characterizing humans on riemannian manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 35(8):1972–1984, 2013. 18

[62] Khai N. Tran, Apurva Bedagkar-Gala, Ioannis A. Kakadiaris, and Shishir K. Shah. Social cues in group formation and local interactions for collective activity analysis. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, volume 1, pages 539–548, 2013. 2, 4, 5, 10, 19, 23, 25

[63] Alessandro Vinciarelli, Maja Pantic, and Hervé Bourlard. Social signal processing: Survey of an emerging domain. *Image and Vision Computing*, 27(12):1743–1759, 2009. 3

[64] Ning Xu, Narendra Ahuja, and Ravi Bansal. Object segmentation using graph cuts based active contours. *Computer Vision and Image Understanding*, 107(3):210–224, 2007. 14

[65] Mohammad Abu Yousuf, Yoshinori Kobayashi, Yoshinori Kuno, Akiko Yamazaki, and Keiichi Yamazaki. Development of a mobile museum guide robot that can configure spatial formation with visitors. In *Intelligent Computing Technology*, pages 423–432. Springer, 2012. 4, 9

[66] Dong Zhang, Daniel Gatica-Perez, Samy Bengio, Iain McCowan, and Guillaume Lathoud. Modeling individual and group actions in meetings: a two-layer hmm framework. In *IEEE International Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, pages 117–124, 2004. 2